

Comments on Jonathan Cohen's, *The Red and the Real*

Andy Egan

Rutgers University & Arché Philosophical Research Centre, University of St.

Andrews

andyegan@rci.rutgers.edu

Jonathan Cohen has written a very rich, interesting, and thought-provoking book, which brings together and pushes forward a line of research that he's been working on, and which I've been profiting from reading, for quite some time. I very much enjoyed reading it, and I've learned a lot from it. I won't voice any very deep disagreements with Cohen in these comments. I count myself as an ally of Cohen's in most of the disputes he's engaging in here. Or at least, I count us as members of camps that are only going to fall out very late in the revolution.

I'll start off by laying a little bit of groundwork for people who haven't read the book, or haven't read it recently. I'll say a little bit about the thesis – *relationalism* about the colors – that Cohen is advocating in the book, and a bit about his central line of argument for it, which is based on facts about perceptual variation within and between observers. I'll bring up a few quibbles and ask a few clarificatory questions along the way.

I'll then pitch an alternative view of the colors, which tries to handle the same phenomena of perceptual variation by taking the contents of color experience to be *de se*. (I'm not sure whether I endorse this view – I waver between it and some kind of dispositionalism that's probably equivalent to Cohen's favored role functionalist view – but I figure it might be useful to complicate the taxonomy by putting it on the table, and to see how it compares to the kind of view that Cohen is advocating.)

2. Cohen's big project: Relationalism.

Cohen's central project in the book is to advocate a fairly broad view-type, rather than a single super-specific view about the nature of the colors. The view-type he's advocating is *Relationalism*, and in order to be a relationalist, you have to think that an object's having a given color is, somehow or other, intimately bound up with its standing in certain relations to subjects.

So on this kind of view, there's no such property as *being-just-plain-red*. Instead, there are a bunch of properties – really a *lot* of properties – of the type, *being red for S in C*. Those are the properties that are represented in color perception, and they're the properties attributed to things when we say e.g. "that rose is red".

I just said that, on a relationalist view, an object's having a given color is "intimately bound up with its standing in certain relations to subjects." It's worth clarifying that a bit – in order to be a relationalist, you have to think that an object's being colored is bound up its relations to subjects in a *constitutive*, rather than a *reference-fixing*, way. It's not just that the reference of color terms is fixed by relations to speakers, perceivers, or whatever/whoever. It's that the properties denoted by color terms are relational properties – properties of standing in a certain sort of relation to observers of a certain type.

(Cautionary note: The preceding presentation is slightly suboptimal, since there's a little bit of a complication in Cohen's story, where the color properties represented in perception aren't the same ones that are the referents of color terms. More on this in a minute.)

This brings us to my first clarificatory question: One locution Cohen uses a lot is “constituted in terms of” – he often says, for example, that the colors “are constituted in terms of” their relations to subjects. Is it fair to gloss the claim that colors are constituted in terms of relations to subjects as, “they’re relational properties”, as I just did? Or by glossing it that way, am I leaving something out, or running over a distinction?

For most of the book, Cohen is carefully neutral about just what species of relationalism we ought to go in for, and is interested just in selling his reader on the virtues of going for *some* kind of relationalist view over the non-relationalist alternatives. Toward the end of the book, he comes out in favor of a *role-functional* view of colors – to be red for S in C, for example, is to have some property or other that disposes things to look red to S in C. (There are, and Cohen discusses, various complications about “looks red”, and about using “looks red” in an analysis of “is red”. But since I don’t have anything new to say about these kinds of worries, I’ll just flag them and move on.)

One thing that’s useful about the way Cohen structures the book, leading with a lot of discussion of relationalism generally, and holding the discussion of his favored version of relationalism until the end, is that it helpfully focuses the argumentative attention, for most of the book, on the virtues and potential vices of relationalism, rather than its particular implementations. I’ll follow his lead on this and not spend too much time on the particular version of representationalism Cohen winds up going for.

3. The argument(s) from perceptual variation

The central line of argument for relationalism is based on variation in color perception, and the implausibility of picking a unique perceiver, or circumstance, in which the colors are being represented veridically.

Very quickly:

1. We see a lot of variation in color perception between observers, and within a single observer in different perceptual circumstances. (The shirt looks red to me and green to you, or red to me and orange to the martian, or red to me and who-knows-how with respect to color to the tetrachromat, or red to me in sunlight but orange to me in fluorescent lighting, or...)
2. In (at least some of) these cases of perceptual variation, it's implausible to say that there's a unique privileged observer, or a unique privileged perceptual circumstance, by/in which the one true color of the object is getting veridically represented.
3. It's better to avoid having to try to find some grounds for privileging one observer or circumstance over the rest, and to say instead that everybody's getting it right.
4. The best way to do that is to go for a relationalist view. If we're relationalists, we can say that S1's experience in C1 represents the shirt as red-for-S1-in-C1, while S2's experience in C2 represents it as orange-for-S2-in-C2, etc. And then there's no need to accuse anybody of misrepresentation.

I endorse this line of argument wholeheartedly. Or rather, I endorse the first bit – the bit that happens in (1-3) – wholeheartedly, and I’m sympathetic to, though not wholeheartedly committed to, the positive proposal. So I won’t say much about this, since (a) I pretty much endorse this line of argument, and (b) I’m sure that plenty of other people will have critical things to say about it.

4. A worry: color proliferation, and loss of common subject matter.

Here’s something to be worried about at this point: Suppose that we’ve got all of these subject- and circumstance-relativised colors. My perceptual experiences in C1 are going to trade in Egan-in-C1 colors, while your perceptual experiences in C2 trade in you-in-C2 colors. Then if each of our color *talk* trades in the same colors as our color *perception*, we’re going to wind up talking past each other, and it’s going to be somewhere between impossible and incredibly unlikely that any two people are ever going to be talking about the same subject matter when they’re talking about the colors of things. That seems bad.

Here is Cohen’s fix for this potential problem: The language picks up on *coarse-grained* relational properties – properties like *red for Ks in C-type circumstances*, rather than *red for S1 in super-specific-circumstance C1*. This is a very interesting and ingenious move.

It gets us – at least plausibly, in very many cases – a common subject matter for people to be talking about in their discourse about colors. Plausibly, it will

happen quite a lot that, when people who are talking to each other, their color talk will pick up on the same (fairly general) perceiver-types and circumstance-types.

It also gives us the resources to explain perceptual error. Another thing you might reasonably be worried about when thinking about the proliferation of colors in Cohen's picture is that it's going to turn out that S's perceptual experiences in C couldn't misrepresent the S-in-C-relativised colors of things. And if those are the only color properties that S's visual perception in C trades in, then it's hard to see how color misrepresentation is going to be possible. But color misrepresentation clearly *is* possible. The appeal to coarse-grained color properties gives Cohen the resources to explain this: Often, when you attribute colors to things on the basis of perception, you go beyond what's right there in the perceptual experience – the fine-grained property, which you correctly represent the object as having – to some broader, coarse-grained property, which it could turn out that the object *doesn't* have, since you were looking at the thing in non-standard circumstances (or, more weirdly, because you're not a member of the perceiver-kind that you thought you were a member of). The circumstances we care about, and which our coarse-grained-color attributions are relativised to, aren't necessarily the circumstances we're actually in.

This is, I think, a very nice story.

I have one lingering worry, which is that this story doesn't deliver a theoretically desirable sort of *contrariness* in perceptual content between spectrum inverted subjects, people who identify different color chips as *unique blue*, etc..

One might have thought (and I would like to be able to think) that when Ernie and Vert are spectrum inverted relative to each other, or when Ron and Hermione differ about whether a given object is unique blue or instead is a bit greenish, there is a representational incompatibility, not just in language – not just in their verbalized color-attributions – but also in perception.

It's natural to think – and I would like to be able to think – that the people who are spectrum inverted (or spectrum-shifted, or...) with respect to each other are representing, not just differently, but *conflictingly*, or *incompatibly*.

Relationalism as Cohen characterizes it doesn't deliver this.

(I weasel with “as Cohen characterizes it” because I'm not confident about whether the sort of view that I'm about to characterize should be counted as a sort of relationalism or not.)

5. Another alternative: They're centering features.

I'm about to make a case for the view that the contents of color perception are *de se*, and that the colors are what I've called elsewhere {Egan:2006tw, Egan:2006uq, Egan:2010wc} *centering features*. I'll start off with a quick introduction to *de se* content, with some motivation but without any real argument for the framework:

Distinguish two different ways of thinking about the objects of propositional attitudes, both of which are instances of possibility-carving views – views that think of representational content as things that impose a division on a space of alternatives. (For advocacy of such views, see for example {Stalnaker:1978tr,

Stalnaker:1984ti}).) For simplicity, let's think just about the versions of these kinds of views that *identify* propositions with regions of a possibility space.

On one kind of picture, the points in the possibility space are possible worlds. On this kind of picture, it's natural to think about the features attributed to objects in, for example, perception as functions from possible worlds to extensions or, equivalently, as functions from objects to possible-worlds propositions.

On another kind of picture (the kind advocated, for example, by David Lewis in "Attitudes De Dicto and De Se" {Lewis:1979vk}), the points in the possibility space are possible *individuals* (or time-slices of individuals), or *centered worlds*. The motivation for adopting this second kind of picture is to model the sorts of phenomena that Lewis talked about, and that Perry talked about in "The Problem of the Essential Indexical" {Perry:1979vc}: identifying something for Perry to come to believe when he realizes that he's the messy shopper, something for Lewis's two gods to be ignorant of, something that captures the doxastic similarity between all of the people who think that Gilmore Lake is before them, or (to import an example from {Kaplan:1989te}) all of the people who think that their pants are on fire, by identifying something that they all believe.

(In the burning pants case, for example, what's in common to all the people who are jumping around, shrieking, looking for fire extinguishers, and saying "my pants are on fire", is that they all believe the centered worlds proposition that includes all and only those $\langle w, t, i \rangle$ triples such that i 's pants are on fire at t in w .)

On this kind of picture, we can think of the features attributed to objects in (for example) perception as, once again, functions from points in the possibility

space to extensions. (Or, equivalently, as functions from objects to regions of the possibility space.) But now, since the points in our possibility space are centered worlds rather than worlds, they'll be functions from centered worlds to extensions, or from objects to centered worlds propositions (i.e., sets of centered worlds). Call these objects *centering features*.

There's an important distinction to draw between two kinds of centered-worlds propositions, and between two different kinds of centering features. Among centered-worlds propositions, there's an important difference between the *boring* centered-worlds propositions, which don't distinguish between positions within a world – for any given world w , they either include every centered world in w or none of them – and the *interesting* centered-worlds propositions, which sometimes draw distinctions between different positions within a given world.

There's an analogous distinction to draw between centering features. It's easiest to state if we think about the centering features as functions from objects to centered-worlds propositions. The *boring* centering features are the ones that always deliver a boring centered-worlds proposition, and the interesting ones are the ones that sometimes deliver an interesting centered-worlds proposition. (Equivalently: the boring centering features are the ones that always deliver the same extension for any two centered worlds that share a world component, and the interesting ones are the ones that aren't boring.)

Apologies for the long setup. Here's the alternative proposal that I think probably counts as a version of relationalism, though I'm not quite sure, and I'm curious to hear what Cohen thinks: Colors are interesting centering features.¹

So for example, we might say that attributing *being green* to Kermit delivers the centered worlds proposition that is true in $\langle w, t, i \rangle$ iff Kermit is, at t , disposed to look green to i in w .

Alternatively, we could go with a version that's more directly analogous to Cohen's view (and better able to handle intrapersonal perceptual variation): Attributing *being green* to Kermit delivers the centered worlds proposition that's true in $\langle w, t, i \rangle$ iff Kermit is disposed to look green to i in the circumstances i occupies at t in w .

This kind of view has some nice features, as compared to the kinds of relationalist views that Cohen considers:

First, it delivers, like the views Cohen considers, perceptual variation without error. Because the contents of color attributions are interesting centered worlds propositions, it can turn out that I correctly attribute e.g. *being unique green* to something that you correctly attribute *being bluish green* to.

Second, and unlike the views Cohen considers, we don't have to shift to coarser-grained color properties in order to get a common subject matter for everybody's color-talk. There's just one color *green*, just one color *red*, etc. (Or anyway, just one color *emerald*₃₂, just one color *scarlet*₂₆ – we'll obviously still want

¹ I've tentatively half-endorsed this kind of view in {Egan:2006uq} and {Egan:2010wc}. I deny it in {Egan:2006tw}, where I argue that we should think not of the colors, but of the non-color *appearance properties* that Shoemaker introduces {Shoemaker:1994tj, Shoemaker:2000vf, Shoemaker:2001vt, Shoemaker:2003vj, Shoemaker:2006uq},¹ as centering features. I'm still conflicted about it.

to allow for the existence of lots of precise shades of the determinable colors like *red* and *green*.) There's no need for a proliferation of *red to S in C* properties. The accounts Cohen considers square the perceptual variation with the absence of error by making it so that different people's color experiences are trading in different properties, all of which the object really has. That, obviously, forces him to allow for a serious proliferation of color properties. The centering feature account, on the other hand, squares the perceptual variation with the absence of error not by multiplying contents, but by attributing to people's color experiences a kind of content that can vary in truth-value across perceivers. So there's no need to multiply color properties.

Finally, I think the biggest benefit of a centering-feature story is that it delivers contrariness: When Vert is spectrum inverted relative to Ernie, Ernie's and Vert's visual experiences of Kermit have incompatible contents. Not incompatible in the sense that they can't both be veridical, but incompatible in the sense that one couldn't veridically represent Kermit as (simultaneously) having both the color that Ernie's experience attributes to him and the color that Vert's experience attributes to him. (Since, for example, *being green* always delivers an extension that's disjoint from the one *being red* delivers, for a given centered world.)

Which leads to my second clarificatory question: To what extent does this count as a relationalist view, and to what extent am I in Cohen's camp?

On the one hand, the views are very similar: on both kinds of views, the colors of things are intimately tied up with their relations to perceivers.

On the other hand, this kind of view is importantly different from the kinds of views Cohen talks about. On the kinds of relationalist views Cohen discusses, there's a proliferation of distinct subject-indexed and circumstance-indexed color properties. On the centering feature view, it could be that the relationality doesn't give rise to any proliferation.

(Though there might still be circumstance-based proliferation – for example, maybe we want something like the *in-normal-circumstances* colors and the *in-present-circumstances* colors in order to deal with color constancy phenomena, where parts of an in-some-sense-uniformly colored surface look in-some-sense-differently colored on account of differences in illumination.)

I'm not sure which of these is more central to how Cohen's thinking about relationalism, or whether there's anything more than a terminological/taxonomical issue here; I look forward to learning what he thinks about this.